UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/812,773	03/30/2004	Paul F. Russell	BEA920030023US1	4984

49474 7590 08/09/2006

LAW OFFICES OF MICHAEL DRYJA
704 228TH AVE NE
#694
SAMMAMISH, WA 98074

EXAMINER	
SONG, JASMINE	
ART UNIT	**PAPER NUMBER**
2188	

DATE MAILED: 08/09/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)
	10/812,773	RUSSELL ET AL.
	Examiner	Art Unit
	Jasmine Song	2188

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *30 March 2004*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-18* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☒ Claim(s) *8 and 9* is/are allowed.

6)☒ Claim(s) *1,2,4-7 and 10-18* is/are rejected.

7)☒ Claim(s) *3* is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☒ The drawing(s) filed on *30 March 2004* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)

2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)

3) ☒ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08) Paper No(s)/Mail Date *03/30/2004*.

4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .

5) ☐ Notice of Informal Patent Application (PTO-152)

6) ☐ Other: _____.

Detailed Action

Specification

1. The lengthy specification has not been checked to the extent necessary to determine the presence of all possible minor errors. Applicant's cooperation is requested in correcting any errors of which applicant may become aware in the specification.

Drawings

2. The drawings filed on 03/30/2004 have been approved by the Examiner.

Oath/Declaration

3. The applicant's oath/declaration has been reviewed by the examiner and is found to conform to the requirements prescribed in 37 C.F.R. 1.63.

Information Disclosure Statement

4. The information disclosure statement (IDS) submitted on 03/30/2004 is in compliance with the provisions of 37 CFR 1.97. Accordingly, the information disclosure statement is being considered by the examiner.

Claim Rejections - 35 USC § 103

5. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

6. Claims 1-2,4-7 and 10-12,14-18 are rejected under 35 U.S.C. 103(a) as being unpatentable over Mathews et al., US 2002/0133678 A1, in view of Sandstrom., US 2001/0013090 A1.

Regarding claim 1, Mathews teaches a method comprising:

A plurality of computing units (it is taught as a plurality of processors 102-106 as shown in Fig.1), for each computing unit of a plurality of computing units, allocating a portion of the logically contiguous section of memory addressable (it is taught allocating the effective address of the processor's mpdata region, section 0043 and section 0042, line 7) by a pointer (it is taught as a handle value, section 0029, line 3-4) plus a static offset corresponding to the computing unit (it is taught as the address of the beginning of each processor's section which is a fixed and per processor value, section 0043), wherein the static offset for each computing unit is equal to a static offset initially determined at initial allocation of memory for the plurality of computing units (it is taught as the offset for each processors is constant and it is the address of the beginning of each processor's section in a fixed location and it is initially determined at initial allocation of memory, section 0012, last three lines).

Mathews also mentions allocating a limited amount of virtual space due to the limited address space and allocation of free storage as needed, storage is allocated using a first fit algorithm which possibly but not necessarily maps non-contiguous physical storage (section 0032-0034), Mathews does not clearly and specifically teach mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory.

However, Sandstrom specifically teaches mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory (section 0012, lines 5-8).

It would have been obvious to the ordinary skill in the art at the time the invention was made to utilize the teachings of Sandstrom into Mathews' system such as mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory because this will release any of the physically non-contiguous memory segment when or if it is found not be in use by any of the program processes and make it available for reuse (section 0012, last four lines), therefore, optimizing the use memory when it is needed by the processors.

Accordingly, one of ordinary skill in the art would have recognized this and concluded that they are from the same field of endeavor. This would have motivated one of ordinary skill in the art to implement the above combination for the advantages set forth above.

Regarding claim 2, Mathews teaches the portion of the logically contiguous section of memory allocated for each computing unit includes memory local to the computing unit (it is taught as the per processor sections are placed in memory that is local to the processors; section 0042, lines 7-10 and section 0027, lines 5).

Regarding claim 4, Mathews teaches further comprising determining the static offset for each computing unit as equal to the static offset initially determined at the initial allocation of the memory for the plurality of computing units (it is taught as offset is the address of the beginning of each processor's section in a fixed location, offset is determined while the storage areas are allocated initially when the system is initialized; section 0012, last three lines).

Regarding claim 5, Mathews teaches further comprising at the initial allocation of the memory for the plurality of computing units:

determining the static offset for each computing unit of the plurality of computing units (it is taught as offset is the address of the beginning of each processor's section in a fixed location, offset is determined while the storage areas are allocated initially when the system is initialized; section 0012, last three lines); and,

for each computing unit of the plurality of computing units, allocating a portion of memory addressable (it is taught allocating the effective address of the processor's mpdata region, section 0043 and section 0042, lines 7) by a pointer (it is taught as a handle value, section 0029, line 3-4) plus a static offset corresponding to the computing

unit (it is taught as the address of the beginning of each processor's section which is a fixed and per processor value).

Regarding claim 6, Mathews teaches further comprising dynamically passing out the portion of the memory to each computing unit of the plurality of computing units as the computing units need additional memory (it is taught as the storage areas are allocated on an as needed basis, section 0012, last three lines and section 0032, lines 5-7).

Regarding claim 7, Mathews teaches the computing unit is one of a computing node and a processor (it is taught as the computing unit is one of a processor as shown in Fig.1).

Regarding claim 10, Mathews teaches a system comprising:

a plurality of computing nodes (it is taught as a plurality of processors as shown in Fig.1);

memory shared by the plurality of computing nodes (section 0020, it is taught as a plurality of processors share a system memory); and,

for each computing node of a plurality of computing nodes, allocating a portion of the logically contiguous section of memory addressable (it is taught allocating the effective address of the processor's mpdata region, section 0043 and section 0042, line 7) by a pointer (it is taught as a handle value, section 0029, line 3-4) plus a static offset

corresponding to the computing node (it is taught as the address of the beginning of each processor's section which is a fixed and per processor value, section 0043), wherein the static offset for each computing node is equal to a static offset initially determined for allocating a portion of memory (it is taught as the offset for each processors is constant and it is the address of the beginning of each processor's section in a fixed location and it is initially determined at initial allocation of memory, section 0012, last three lines).

Mathews also mentions allocating a limited amount of virtual space due to the limited address space and allocation of free storage as needed, storage is allocated using a first fit algorithm which possibly but not necessarily maps non-contiguous physical storage (section 0032-0034), and he also teaches an allocating mechanism (it is taught as the meaning of allocating, section 0032-0034), Mathews does not clearly and specifically teach mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory.

However, Sandstrom specifically teaches mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory (section 0012, lines 5-8).

It would have been obvious to the ordinary skill in the art at the time the invention was made to utilize the teachings of Sandstrom into Mathews' system such as mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory because this will release any of the physically non-contiguous memory segment when or if it is found not be in use by any of the program processes

and make it available for reuse (section 0012, last four lines), therefore, optimizing the use memory when it is needed by the processors.

Accordingly, one of ordinary skill in the art would have recognized this and concluded that they are from the same field of endeavor. This would have motivated one of ordinary skill in the art to implement the above combination for the advantages set forth above.

Regarding claim 11, Mathews teaches the portion of the logically contiguous section of memory allocated to each computing node includes memory local to the computing node (it is taught as the per processor sections are placed in memory that is local to the processors; section 0042, lines 7-10 and section 0027, lines 5).

Regarding claim 12, Mathews teaches each of the plurality of computing nodes comprises a single processor (Fig.1, it is taught as each of the plurality of computing nodes is a single processor).

Regarding claim 14, Mathews teaches a system comprising:

a plurality of computing nodes (it is taught as a plurality of processors as shown in Fig.1);

memory shared by the plurality of computing nodes (section 0020, it is taught as a plurality of processors share a system memory); and,

for each computing node of a plurality of computing nodes, allocating a portion of the logically contiguous section of memory addressable (it is taught allocating the effective address of the processor's mpdata region, section 0043 and section 0042, line 7) by a pointer (it is taught as a handle value, section 0029, line 3-4) plus a static offset corresponding to the computing node (it is taught as the address of the beginning of each processor's section which is a fixed and per processor value, section 0043), wherein the static offset for each computing node is equal to a static offset initially determined for allocating a portion of memory (it is taught as the offset for each processors is constant and it is the address of the beginning of each processor's section in a fixed location and it is initially determined at initial allocation of memory, section 0012, last three lines).

Mathews also mentions allocating a limited amount of virtual space due to the limited address space and allocation of free storage as needed, storage is allocated using a first fit algorithm which possibly but not necessarily maps non-contiguous physical storage (section 0032-0034), Mathews does not clearly and specifically teach mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory.

However, Sandstrom specifically teaches mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory (section 0012, lines 5-8).

It would have been obvious to the ordinary skill in the art at the time the invention was made to utilize the teachings of Sandstrom into Mathews' system such as mapping

a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory because this will release any of the physically non-contiguous memory segment when or if it is found not be in use by any of the program processes and make it available for reuse (section 0012, last four lines), therefore, optimizing the use memory when it is needed by the processors.

Accordingly, one of ordinary skill in the art would have recognized this and concluded that they are from the same field of endeavor. This would have motivated one of ordinary skill in the art to implement the above combination for the advantages set forth above.

Regarding claim 15, Mathews teaches a computing node comprising:

a plurality of processors (Fig.1, processors 102-106);

memory shared by the plurality of processors (section 0020, it is taught as a plurality of processors share a system memory); and,

for each computing processor of a plurality of processors, allocating a portion of the logically contiguous section of memory addressable (it is taught allocating the effective address of the processor's mpdata region, section 0043 and section 0042, line 7) by a pointer (it is taught as a handle value, section 0029, line 3-4) plus a static offset corresponding to the processor (it is taught as the address of the beginning of each processor's section which is a fixed and per processor value, section 0043), wherein the static offset for each processor is equal to a static offset initially determined for allocating a portion of memory (it is taught as the offset for each processors is constant

and it is the address of the beginning of each processor's section in a fixed location and it is initially determined at initial allocation of memory, section 0012, last three lines).

Mathews also mentions allocating a limited amount of virtual space due to the limited address space and allocation of free storage as needed, storage is allocated using a first fit algorithm which possibly but not necessarily maps non-contiguous physical storage (section 0032-0034), and he also teaches an allocating mechanism (it is taught as the meaning of allocating, section 0032-0034), Mathews does not clearly and specifically teach mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory.

However, Sandstrom specifically teaches mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory (section 0012, lines 5-8).

It would have been obvious to the ordinary skill in the art at the time the invention was made to utilize the teachings of Sandstrom into Mathews' system such as mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory because this will release any of the physically non-contiguous memory segment when or if it is found not be in use by any of the program processes and make it available for reuse (section 0012, last four lines), therefore, optimizing the use memory when it is needed by the processors.

Accordingly, one of ordinary skill in the art would have recognized this and concluded that they are from the same field of endeavor. This would have motivated one of ordinary skill in the art to implement the above combination for the advantages

set forth above.

Regarding claim 16, Mathews teaches the portion of the logically contiguous section of memory allocated to each processor is local to the processor (it is taught as the per processor sections are placed in memory that is local to the processors; section 0042, lines 7-10 and section 0027, lines 5).

Regarding claim 17, Mathews teaches a computing node comprising:

a plurality of processors (Fig.1, processors 102-106);

memory shared by the plurality of processors (section 0020, it is taught as a plurality of processors share a system memory); and,

for each processor of a plurality of processors, allocating a portion of the logically contiguous section of memory addressable (it is taught allocating the effective address of the processor's mpdata region, section 0043 and section 0042, lines 7) by a pointer (it is taught as a handle value, section 0029, line 3-4) plus a static offset corresponding to the processor (it is taught as the address of the beginning of each processor's section which is a fixed and per processor value, section 0043), wherein the static offset for each processor is equal to a static offset initially determined for allocating a portion of memory (it is taught as the offset for each processors is constant and it is the address of the beginning of each processor's section in a fixed location and it is initially determined at initial allocation of memory, section 0012, last three lines).

Mathews also mentions allocating a limited amount of virtual space due to the limited address space and allocation of free storage as needed, storage is allocated using a first fit algorithm which possibly but not necessarily maps non-contiguous physical storage (section 0032-0034), Mathews does not clearly and specifically teach mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory.

However, Sandstrom specifically teaches mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory (section 0012, lines 5-8).

It would have been obvious to the ordinary skill in the art at the time the invention was made to utilize the teachings of Sandstrom into Mathews' system such as mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory because this will release any of the physically non-contiguous memory segment when or if it is found not be in use by any of the program processes and make it available for reuse (section 0012, last four lines), therefore, optimizing the use memory when it is needed by the processors.

Accordingly, one of ordinary skill in the art would have recognized this and concluded that they are from the same field of endeavor. This would have motivated one of ordinary skill in the art to implement the above combination for the advantages set forth above.

Regarding claim 18, Mathews teaches an article of manufacture comprising:

a computer-readable medium (it is taught as a system memory shared by a plurality of processors); and,

for each computing unit of a plurality of computing units, allocating a portion of the logically contiguous section of memory addressable (it is taught allocating the effective address of the processor's mpdata region, section 0043 and section 0042, lines 7) by a pointer (it is taught as a handle value, section 0029, line 3-4) plus a static offset corresponding to the computing unit (it is taught as the address of the beginning of each processor's section which is a fixed and per processor value, section 0043), wherein the static offset for each computing unit is equal to a static offset initially determined at boot time of the plurality of computing units for allocating a portion of memory to each computing units (it is taught as the offset for each processors is constant and it is the address of the beginning of each processor's section in a fixed location and it is initially determined at initial allocation of memory, section 0012, last three lines).

Mathews also mentions allocating a limited amount of virtual space due to the limited address space and allocation of free storage as needed, storage is allocated using a first fit algorithm which possibly but not necessarily maps non-contiguous physical storage (section 0032-0034), Mathews does not clearly and specifically teach mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory.

However, Sandstrom specifically teaches mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory (section 0012, lines 5-8).

It would have been obvious to the ordinary skill in the art at the time the invention was made to utilize the teachings of Sandstrom into Mathews' system such as mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory because this will release any of the physically non-contiguous memory segment when or if it is found not be in use by any of the program processes and make it available for reuse (section 0012, last four lines), therefore, optimizing the use memory when it is needed by the processors.

Accordingly, one of ordinary skill in the art would have recognized this and concluded that they are from the same field of endeavor. This would have motivated one of ordinary skill in the art to implement the above combination for the advantages set forth above.

7. Claim 13 is rejected under 35 U.S.C. 103(a) as being unpatentable over Mathews et al., US 2002/0133678 A1, in view of Venkatsubramanian et al., US 6,886,031 B2.

Regarding claim 13, Mathews teaches the claimed invention as shown in claim 10, Mathew does not teach that each of the plurality of computing nodes comprises a plurality of processors. However, Venkatsubramanian teaches each of the plurality of computing nodes comprises a plurality of processors (Fig.4, Venkatsubramanian

teaches a plurality of SMP systems, each SMP system is referred to as an SMP node and each SMP node comprises a plurality of processors, col.1, lines 36-40).

It would have been obvious to the ordinary skill in the art at the time the invention was made to utilize the teachings of Venkatsubramanian into Mathews' system such as each of the plurality of computing nodes comprises a plurality of processors because this is a multiprocessing SMP architecture and SMP architecture provides fast performance by making different processors available to a number of processes on the machine simultaneously (col.1, lines 26-40).

Accordingly, one of ordinary skill in the art would have recognized this and concluded that they are from the same field of endeavor. This would have motivated one of ordinary skill in the art to implement the above combination for the advantages set forth above.

Allowable Subject Matter

8. Claims 8-9 are allowed.

9. Claim 3 is objected to as being dependent upon a rejected base claim, but would be allowable if rewritten in independent form including all of the limitations of the base claim and any intervening claims.

The following is a statement of reasons for the indication of allowable subject matter: The prior art of record also does not teach or suggest that Upon the logically contiguous section of memory being completely used by the computing units, a number of physically non-contiguous sections of memory are mapped into another logically

contiguous section of memory as claimed in claims 3 and 8 in combination with the other elements set forth in the claimed invention. Claim 9 is dependent on claim 8 and incorporated with the additional allowable subject matters.

Conclusion

10. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

Hansen et al	US 2005/0132250 A1
Sexton et al	US 6499095 B1

11. When responding to the office action, Applicant is advised to clearly point out the patentable novelty which he or she thinks the claims present in view of the state of the art disclosed by the references cited or the objections made. He or she must also show how the amendments avoid such references or objections. See 37 C.F.R. 1.111 (c).

12. When responding to the office action, Applicants are advised to provide the examiner with the line numbers and page numbers in the application and/or references cited to assist examiner to locate the appropriate paragraphs.

13. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Jasmine Song whose telephone number is 571-272-4213. The examiner can normally be reached on 7:30-5:30 (first Friday off).

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Mano Padmanabhan can be reached on 571-272-4210. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

Jasmine Song
Jasmine Song

Patent Examiner

August 5, 2006